

Dimension-splitting for simplifying diffusion in lattice-gas models

Raissa M. D'Souza*

Bell Laboratories, Murray Hill, NJ 07974

Norman H. Margolus†

Center for Computational Science, Boston University, Boston, Massachusetts 02215

Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Mark A. Smith‡

Department of Molecular and Cellular Biology, Harvard University, Cambridge, Massachusetts 02138

New England Complex Systems Institute, Cambridge, Massachusetts 02138

(Dated: February 8, 2008)

We introduce a simplified technique for incorporating diffusive phenomena into lattice-gas molecular dynamics models. In this method, spatial interactions take place one dimension at a time, with a separate fractional timestep devoted to each dimension, and with all dimensions treated identically. We show that the model resulting from this technique is equivalent to the macroscopic diffusion equation in the appropriate limit.

This technique saves computational resources and reduces the complexity of model design, programming, debugging, simulation and analysis. For example, a reaction-diffusion simulation can be designed and tested as a one-dimensional system, and then directly extended to two or more dimensions. We illustrate the use of this approach in constructing a microscopically reversible model of diffusion-limited aggregation as well as in a model of growth of biological films.

Key words: Dimension-splitting, diffusion, lattice-gas automata, transport theory, alternating direction method, microscopic reversibility, parallel computing, fractional timestep method.

I. INTRODUCTION

In computer modeling of physical systems, there is not always a direct correspondence between simple models and computationally efficient ones. What seems like a direct and intuitive description to a human does not necessarily correspond to an algorithm that performs well on a computer. Simplicity and directness can, however, be valuable assets regardless of whether they result in better performance. Simple models are easier to analyze and to translate into correct computer programs. In addition, complex models composed of well-understood simple component models are easy to design and test. Of course what we would really like is to combine simplicity and performance.

Molecular dynamics—numerically simulating the classical or semiclassical behavior of a collection of atoms or molecules—is an example of a direct technique with wide applicability that is conceptually simple. The technique is often used in situations where a computational solution based on differential equations is infeasible due to the complexity of the physical dynamics or the complexity of the boundary conditions. Often, in molecular dynamics simulations, the microscopic dynamics of the particles is simplified in order to enable more computationally efficient and hence larger-scale simulations, as in the dissipative particle dynamics and lattice-gas approaches[1, 2]. In a lattice-gas automaton (LGA), for example, only a small number of particle velocities are allowed, and particles are confined to the sites of a finite spatial lattice. Steps in which particles concurrently hop between lattice sites alternate with steps in which the particles present at each lattice site interact. This simplification of the microscopic details results in models that combine simplicity and performance, while still capturing the full essence of phenomena of interest.

LGA-like models have long been studied in physics as theoretical models which capture some essence of a physical phenomenon. The classic example is the random walk on a lattice, which has been used as a microscopic model of diffusion[3]. The first true LGA models were introduced for the theoretical study of fluid flow[4]. It was realized only much later that the locality, uniformity and spatial regularity of models based on LGA's make them ideal candidates for large-scale simulation on parallel hardware[5, 6, 7]. Since then, LGA simulations have been used to study a variety of physical phenomena, including fluid dynamics, chemical reactions, and changes in phase[8, 9, 10, 11, 12].

*Electronic address: raissa@bell-labs.com

†Electronic address: nhm@mit.edu

‡Electronic address: smith2@fas.harvard.edu

More generally, lattice dynamics which emulate the locality and uniformity of physical dynamics occupy a borderland between abstract theoretical models and practical algorithms: they are often both[13, 14, 15, 16]. The computer model is the same as the theoretical model, unlike models based on partial differential equations (PDE's), for example, where an intervening stage of numerical analysis is needed. Thus lattice models and related modeling techniques have both a theoretical and a practical aspect, providing simple models that we can directly experiment with, visualize and analyze mathematically.

In this paper, we introduce a simple technique for incorporating diffusive phenomena into LGA and other lattice models. This technique is closely related to the classical random walk model, but differs in important respects. First, it incorporates exclusion: only a limited number of particles can occupy each lattice site. The technique we use in this manuscript to implement diffusion with exclusion was first introduced by Toffoli[13]. A second difference involves randomness: it is explicitly provided by concurrently simulating a different invertible LGA as a source of stochasticity. This approach has also been used before in other parallel LGA models[7, 14]. Finally, perhaps the most important distinction from a classical random walk model is that this technique splits each n -dimensional simulation time-step up into n fractional time-steps, updating only one dimension during each fractional step. This is similar to some dimension-splitting and fractional-time-step techniques used in solving PDE's[17, 18], but here we apply it in the context of a lattice-gas with exact particle conservation. Earlier Cellular Automata (CA) models have used dimension-splitting[19, 20], but not for simulating diffusion, and not as a general modeling technique. For a review of other techniques for simulating diffusion with discrete lattice systems see Ref. [21].

LGA models of complex materials and phenomena can be constructed in a modular fashion by combining simpler models—something which is difficult to do with differential equation based models. Qualitative models are particularly easy to construct in this manner, but realistic quantitative models of physical phenomena can also be constructed. For example, to reproduce known chemical behavior, several subsystems can be combined. Each subsystem would consist of independently diffusing chemical species, and interaction between subsystems (*e.g.* chemical reaction steps) would occur in between independent species-diffusing steps[11, 12]. The diffusing steps can be handled with a dimension-splitting approach, resulting in a very simple treatment of the diffusion component of the model, which requires only a small amount of state information at each lattice site. A smaller amount of state not only allows larger simulations to be achieved with given memory resources, but also makes it easier to couple component subsystems together. The uniform treatment of one dimension at a time provides an additional benefit. Consider, for example, a chemical reaction model of the sort described above. Such a model can be developed and tested as a two-dimensional model and then converted into a three-dimensional model by simply applying the same dynamics to an additional dimension.

In general, LGA models have a different domain of applicability than PDE-based models. Since they are based on a kind of molecular dynamics optimized for exact digital computation on massively parallel hardware, LGA's can be constructed to handle complex interactions and complex boundaries—circumstances poorly suited to PDE-based simulations. For example, running a fluid LGA simulation through an MRI reconstruction of a real porous rock involves no more computational work than running it with any other boundary condition. An example illustrating the power of the LGA approach for complex systems with complicated boundaries is a lattice-gas model of amphiphilic fluid flow in porous media[22].¹

In this paper we first review a version of Toffoli's LGA model of diffusion with exclusion. We follow an analysis due to Chopard and Droz[23, 24] to show explicitly how to start with the microscopic model and recover the appropriate macroscopic diffusion equation in the continuum limit. We then introduce a dimension-splitting technique for use in LGA models that involve diffusion. We show explicitly how to map this dimension-reduced model onto Toffoli's LGA model. We then discuss the origins of the conservation laws present in the model. Finally we illustrate the use of this technique as a component for building two composite LGA models: a microscopically reversible model of diffusion-limited aggregation, and a model of growth of biological films. These two models were developed in two dimensions but have also been run and studied in three dimensions.

II. DIFFUSION

Modeling diffusion is a cornerstone for building general models of a wide range of physical phenomena. Examples include models for reaction-diffusion processes, pattern formation, and heat flow[13, 25, 26]. Models of diffusive transport alone encompass such diverse phenomena as flow through porous media[8], flow of fluids in biological systems[27], and flow in geological processes[8, 25].

¹ Note that in Sec. (II A 1) we introduce the lattice-Boltzmann equation (LBE) for our model and discuss some of the distinctions between the LGA and the LBE approaches.

Two standard approaches to modeling diffusion come from opposite extremes: the continuum limit, where diffusion is modeled by a partial differential equation; and the microscopic limit where diffusion is modeled as a collection of particles undergoing random walks (which is meant to capture the phenomenon of Brownian motion). One-dimensional diffusive phenomena were first described from a macroscopic perspective by Fick[28]. The generalization to multiple dimensions is the well known diffusion equation:

$$\frac{\partial \rho(\vec{x}, t)}{\partial t} = D \nabla^2 \rho(\vec{x}, t). \quad (1)$$

Here $\rho(\vec{x}, t)$ is density as a function of space and time, and D (the coefficient of self-diffusion) describes the diffusion rate.

We are interested in modeling diffusive phenomena beginning at the microscopic scale and recovering Eq. (1) in the appropriate limit. We first consider a lattice-gas automata (LGA) model of particles hopping at random along the sites of a periodic lattice. The particles execute simultaneous random walks while obeying an exclusion principle. This model was introduced by Toffoli, and it was extended and studied extensively by Chopard and Droz. They show that in the limit of the infinite lattice, the microscopic particle model maps directly onto the macroscopic continuum description, Eq. (1). Additionally, they show that the Green-Kubo relation for expressing transport coefficients (such as the diffusion coefficient D) in terms of autocorrelation functions is recovered exactly.

Below, we present the model and review the basic analysis of Chopard and Droz. Then we present a new model which introduces a dimension-splitting technique. Then we show, by a remapping, that both models are equivalent.

A. An LGA model of diffusion with exclusion

Consider a two-dimensional square lattice.² We can think of particles hopping between adjacent sites of the lattice with unit velocity. The particles are initialized on lattice sites so at every integer time, the particles occupy a lattice site. The particles can move in one of four directions: north, west, south, or east—corresponding to the four lattice directions. In the language of lattice-gases, we consider a “transport channel”, $N_i(\vec{x}, t)$, along each lattice direction, \hat{c}_i , at each site \vec{x} at time t , where $i \in \{0, 1, 2, 3\}$, and \hat{c}_i is the unit vector along that lattice direction. The transport channel can either be occupied (with a particle to be moved) or empty: $N_i(\vec{x}, t) = 1$ or 0 respectively. As there are four channels at each site, there can be up to four particles at each site.

The dynamics can be decomposed into two portions: the “interaction” phase and the “streaming” phase. We denote the state *after* interaction but *before* streaming as $N'_i(\vec{x}, t)$. Note, the interval for the completion of both phases is considered the unit interval of time. During the streaming phase the particles are transported along the lattice. If there were no interaction phase, the particles would stream along the lattice ballistically, with their velocities never changing.

To implement simultaneous random walks, during the interaction phase the states of the transport channels at each lattice site are randomly permuted.³ The random permutation is applied to all lattice sites simultaneously, with each permuted independently of all other sites. Since the system is initialized with at most one particle in each channel, this restriction is preserved by the dynamics. Thus all the particles execute simultaneous random walks with exclusion. In the model considered by Chopard and Droz, each channel at site \vec{x} at time t was simultaneously permuted by a rotation of $r(\vec{x}, t) \cdot \pi/2$ radians, where $r(\vec{x}, t) \in \{0, 1, 2, 3\}$ is a random number with values that occur with a frequency corresponding to a specified probability distribution. By choosing this probability distribution appropriately, we can control the mean-free path (i.e., the average number of sites a particle advances before having its velocity permuted). Thus after the permutation,

$$N'_i(\vec{x}, t) = N_{i+r(\vec{x}, t)}(\vec{x}, t), \quad (2)$$

where here and in what follows a sum in a subscript, such as $(i + r(\vec{x}, t))$, is taken modulo four. During the subsequent streaming phase the particles are transported along the lattice:

$$N_i(\vec{x}, t + 1) = N'_i(\vec{x} - \hat{c}_i, t) = N_{i+r(\vec{x}-\hat{c}_i, t)}(\vec{x} - \hat{c}_i, t). \quad (3)$$

² A similar argument would apply to models defined on the two-dimensional triangular lattice, with the transport split into a sequence of three operations performed one principle lattice direction at a time.

³ We should note that this approach can be considered a field-theoretic molecular dynamics: the fields, not the particles, are permuted. See Ref. [13] for an early discussion of this distinction.




FIG. 1: A realization of diffusion simulated by a simple CA rule which conserves particle number (mass) but not momentum. Particles move with unit velocity along the lattice directions of a two-dimensional square lattice, of size $L \times L$ where $L = 512$. For each lattice site, at each update, the velocities of the particles (1's) and holes (0's) at that site are randomly permuted. The random permutation used at each lattice site is chosen independently. Starting from an initial uniform block of particles, as shown in (a), we observe the diffusive behavior shown in (b), which corresponds to $t = 360$ complete updates of the lattice.

Though Chopard and Droz considered a more general model, we restrict ourselves to the case where $r(\vec{x}, t)$ is uniformly and independently distributed amongst the four possible values for each (\vec{x}, t) . Figure 1 shows the qualitative diffusive behavior of the model. The system is initialized with a uniform block of particles occupying an otherwise vacant two-dimensional lattice, as shown in Fig. 1(a). After $t = 360$ updates of the full lattice we observe the configuration shown in Fig. 1(b).

1. Establishing the lattice Boltzmann equation

Up to here, we have described the dynamics for a single realization of the lattice, as it is actually implemented and run on a computer in practice. However, for purposes of understanding how to use this technique as a realistic generator of diffusion, we would like to know “what happens on average in the limit” (where these terms have not yet been defined). Figure 1 shows what happens for a particular realization; in this section, we take the average; and in the next section, we take the limit.

We first consider the mathematical idealization of an ensemble of all such simulations. We consider that each member of the ensemble is initialized in a consistent way (as discussed in detail at the end of this section), and evolves independently. There are many quantities one could average over this ensemble to test the model, but the main quantity of interest is the state of each channel: $N_i(\vec{x}, t)$. Taking the expectation of Eq. 3 gives:

$$\langle N_i(\vec{x}, t+1) \rangle = \langle N_{i+r(\vec{x}-\hat{c}_i, t)}(\vec{x} - \hat{c}_i, t) \rangle = \left\langle \frac{1}{4} \sum_{j=0}^3 N_j(\vec{x} - \hat{c}_i, t) \right\rangle = \frac{1}{4} \sum_{j=0}^3 \langle N_j(\vec{x} - \hat{c}_i, t) \rangle. \quad (4)$$

Note that we used the fact that the expectation of a sum is the sum of expectations, regardless of the correlations between distinct channels. Finally, defining the occupation number as $n_i(\vec{x}, t) \equiv \langle N_i(\vec{x}, t) \rangle$ gives the lattice-Boltzmann equation⁴ for our system,

$$n_i(\vec{x}, t+1) = \frac{1}{4} \sum_{j=0}^3 n_j(\vec{x} - \hat{c}_i, t). \quad (5)$$

⁴ A lattice-Boltzmann equation (LBE) describes a dynamical system spatially arrayed on a lattice but with continuous variables at each site (as opposed to discrete variables, which is the case for CA's). LBE's have many properties which make them better suited for modeling hydrodynamic behavior. Yet in situations where microscopic noise, kinetic fluctuations and interparticle correlations play an essential role, one must use LGA models. For a discussion of the striking difference in phenomenology accessible to LGA as opposed to LBE see Ref. [29].

Averaging over the ensemble in this way transforms our discrete Boolean variables into continuous probabilities.

Equation (5) is one of the few instances where an interesting physical dynamics is exactly described by a linear lattice-Boltzmann equation (as discussed above, correlations between channels do not matter). Typically, there are interactions between channels, so in order to perform the ensemble average (as in Eq. (4)), one must make the approximation that the channels are uncorrelated (*i.e.*, the assumption of molecular chaos), and the resulting equations are nonlinear. In the present case there are no interactions between the channels because the streaming and permutations are *data blind*, meaning they are not conditional on the state of the channels. As there is still confusion about the correspondence between microscopic dynamics and the macroscopic limit, especially in the context of the increase in macroscopic entropy which may accompany a microscopically reversible dynamics[30, 31], models described by linear equations may be illuminating. Exact solutions can be obtained and discrepancies with the continuous limit discussed rigorously[24].

2. Mapping onto the diffusion equation

If we now view each $n_i(\vec{x}, t)$ as a smooth function of a continuous spacetime, we can derive a partial differential equation that governs the dynamics of the total density, $\rho(\vec{x}, t)$, of particles at a given site:

$$\rho(\vec{x}, t) \equiv \sum_{i=0}^3 n_i(\vec{x}, t). \quad (6)$$

Combining this with Eq. (5) we can write out the evolution of the density,

$$\begin{aligned} \rho(\vec{x}, t+1) &= \sum_{i=0}^3 n_i(\vec{x}, t+1) \\ &= \frac{1}{4} \sum_{i=0}^3 \left[\sum_{j=0}^3 n_j(\vec{x} - \hat{c}_i, t) \right] = \frac{1}{4} \sum_{i=0}^3 \rho(\vec{x} - \hat{c}_i, t). \end{aligned} \quad (7)$$

The limit of interest is the diffusive regime where the time step, Δt , and lattice spacing, $\Delta x = |\hat{c}_i|$, approach zero: $\Delta x \rightarrow 0$, and $\Delta t \rightarrow 0$, while $(\Delta x)^2 / \Delta t \rightarrow \text{constant}$. Note in this regime (of diffusion on an infinite lattice), the propagation velocity, proportional to $\Delta x / \Delta t$, becomes infinite. This unphysical propagation velocity is actually a well known pathology of the diffusion equation, Eq. (1).

Taylor expanding the spatial terms in Eq. (7), keeping those terms to order $(\Delta x)^2$, we obtain

$$\begin{aligned} \rho(\vec{x}, t+1) &= \frac{1}{4} \sum_{i=0}^3 \sum_{j=0}^3 \left[n_j(\vec{x}, t) + \Delta x (\hat{c}_i \cdot \nabla) n_j(\vec{x}, t) + \frac{(\Delta x)^2}{2} (\hat{c}_i \cdot \nabla)^2 n_j(\vec{x}, t) \right] \\ &= \rho(\vec{x}, t) + \frac{(\Delta x)}{4} \sum_{i=0}^3 \sum_{j=0}^3 (\hat{c}_i \cdot \nabla) n_j(\vec{x}, t) + \frac{(\Delta x)^2}{8} \sum_{i=0}^3 \sum_{j=0}^3 (\hat{c}_i \cdot \nabla)^2 n_j(\vec{x}, t). \end{aligned} \quad (8)$$

After noting that $(\hat{c}_i = -\hat{c}_{i+2})$ and $[(\hat{c}_i \cdot \nabla)^2 + (\hat{c}_{i+1} \cdot \nabla)^2] = \nabla^2$, we can simplify Eq. (8):

$$\rho(\vec{x}, t+1) = \rho(\vec{x}, t) + \frac{(\Delta x)^2}{4} \nabla^2 \rho(\vec{x}, t). \quad (9)$$

Taylor expanding to order (Δt) we obtain the diffusion equation:

$$\frac{\partial \rho(\vec{x}, t)}{\partial t} = D \nabla^2 \rho(\vec{x}, t), \quad (10)$$

where the diffusion constant, $D = (\Delta x)^2 / 4 \Delta t$.

We have shown above that the ensemble as a whole reproduces the diffusion equation in the limit of an infinitely fine lattice and continuous time. However, for purposes of an actual LGA simulation, all we have is a single member of the ensemble in a discrete spacetime. In particular, the state of each individual channel is a 0 or 1 and not a real number between 0 and 1 as results from a lattice-Boltzmann treatment. So we want to argue that a typical member

of the ensemble—as generated by a single run of a simulation—provides a realistic rendition of the diffusion of a collection of particles.

The main thing we want in this respect is that the expected number of particles in a patch of lattice sites is in some sense the same as the expected number in the corresponding region of continuous space. More specifically, we demand equivalence at the mesoscopic scale (the scale at which patches are large compared to the lattice spacing but small compared to the characteristic scale of the system as a whole). Note that it is also at this mesoscopic scale that we demand equivalence between LGA and lattice-Boltzmann equations. We could sum $\rho(\vec{x}, t)$, as obtained from the lattice-Boltzmann formulation, over a patch of lattice sites and obtain a quantity representing the number of particles in that patch. On the other hand, we could generate individual members of the ensemble from simulations according to the LGA formulation, count the number of particles in the patch, and then take the ensemble average of this count. Both approaches give exactly the same quantity. Much as is the case for the separate channels feeding into a single site, as discussed following Eq. (4), the expectation over the ensemble for this count in the LGA formulation is the same as the sum of the expectations giving $n_i(\vec{x}, t)$ in the lattice-Boltzmann formulation.

This is true as long as each member of the ensemble is chosen to have the appropriate average mesoscopic initial conditions. Equation (5) provides an exact description of any ensemble that starts with the correct average at each lattice site, no matter how strong the correlations between sites. However, not every member of such an initial ensemble is necessarily “typical” in that it could be consistent with the mesoscopic conditions, yet have a very unusual pattern of correlations between the cells; for example, the particles might align to form an image of a face. To prevent this in practice, one initializes each *channel* at a relevant site (\vec{x}, t) independently at random to be a 1 with a probability given by $n_i(\vec{x}, t)$ and to be 0 with a probability given by $1 - n_i(\vec{x}, t)$. This generates a sample of the maximum entropy ensemble that is consistent with the initial specified averages. Moreover, the sample is self-averaging in the sense that the average density of particles in any patch around a point will closely match the initial average at that point. Finally, at any point in the simulation, the configuration is likely to be a typical member of the ensemble because the dynamics only serves to increase coarse-grained entropy, not generate highly correlated patterns. It is these typical ensemble members that provide the most realistic rendition of the diffusion of a collection of particles.

While the above analysis is very promising, we should mention some caveats of the model. In particular, the simulation cannot match more complicated statistics of a real physical system. For example, the fluctuations in the particle number are bounded above by the number of channels in any patch whereas in principle, there is no upper bound in a continuum. Also, the specific microscopic details will never be the same as in a real physical system or from one run to the next. While the exact configuration of all the particles do not matter to the average density of the diffusing species alone, they do matter a great deal when they are coupled to a system that is very sensitive to initial conditions, such as the reversible aggregation model that will be described in Sec. III A, where microscopic fluctuations dictate the dynamics.

B. The dimension-splitting technique

Our approach is to consider the multi-dimensional diffusion process as a sequence of one-dimensional diffusion events. Each one-dimensional diffusion step consists of an interaction (*i.e.*, mixing) phase and a transport phase. We need consider only a single transport channel and its opposite channel: $M_0(\vec{x}, t)$ and $M_1(\vec{x}, t)$. First consider the mixing phase. An unbiased, random binary variable, $\eta(\vec{x}, t) \in \{0, 1\}$, $\langle \eta(\vec{x}, t) \rangle = 1/2$, is sampled to determine whether the states of the two channels at site (\vec{x}, t) are interchanged:

$$\begin{aligned} M'_0(\vec{x}, t) &= [1 - \eta(\vec{x}, t)]M_0(\vec{x}, t) + \eta(\vec{x}, t)M_1(\vec{x}, t) \\ M'_1(\vec{x}, t) &= [1 - \eta(\vec{x}, t)]M_1(\vec{x}, t) + \eta(\vec{x}, t)M_0(\vec{x}, t). \end{aligned} \quad (11)$$

Then the updated channels are transported along the \hat{x} direction, as shown in Fig. 2(a):

$$\begin{aligned} M_0(\vec{x}, t) &= M'_0(\vec{x} - \hat{x}, t - 1) \\ M_1(\vec{x}, t) &= M'_1(\vec{x} + \hat{x}, t - 1). \end{aligned} \quad (12)$$

Once these two substeps are complete, we consider diffusion along the \hat{y} direction which is also composed of the analogous mixing and transport steps (*i.e.*, substitute \hat{y} for \hat{x} in Eq. (12)). As shown in Fig. 2, we first use the two channels to transport data along the horizontal direction, then we use these same two channels to transport data along the vertical direction. We thus reinterpret the function of the same two degrees of freedom.

In this decomposition of a two-dimensional dynamics into two one-dimensional dynamics, instead of four transport channels at each site, we need only two. Thus at each lattice site, three bits of state are required: two bits, $M_\gamma(\vec{x}, t)$ where $\gamma \in \{0, 1\}$, denote the presence or absence of a particle in each of the two channels respectively; the third bit, $\eta(\vec{x}, t)$, is a binary pseudorandom variable. Note, the conventional approach requires six bits of state at each site.




FIG. 2: An illustration of how the degrees of freedom are reinterpreted using a partitioning scheme. A snapshot of a lattice containing a single particle, represented by the shaded circle, is shown. The arrows represent the allowed directions of motion. In (a) the degrees of freedom are considered to be the directions of motion along the horizontal axis. In (b) they are the directions along the vertical axis.

The number of degrees of freedom interacting at any one time is simplified (and remains so regardless of overall dimensionality). Yet we end up executing a sequence of four substeps (mix, then transport along \hat{x} , mix then transport along \hat{y}). In contrast, the full two-dimensional problem requires only two substeps (one mixing, one transport). Note also that extending the dynamics to one more dimension requires only including an additional sequence of mixing followed by transport (*i.e.*, an additional fractional timestep, with the transport along the new dimension).⁵

In the two-dimensional diffusion model considered, we find that at the end of one complete update (*i.e.*, one update along \hat{x} and one update along \hat{y}) a particle has advanced one unit along each dimension—one step along an oblique lattice with lattice spacing $|\hat{s}_i| = \sqrt{2}|\hat{c}_i|$, where $|\hat{c}_i|$ is the original lattice spacing. The mapping onto this oblique lattice is shown in Fig. 3.

1. Mapping onto the diffusion equation

We show that this dimension-reduced approach to diffusion maps onto the Chopard-Droz analysis discussed in Sec. II A, and thus also onto the diffusion equation, Eq. (1), in the appropriate regime. For simplicity we again consider the two-dimensional case, and again deal with the average over an ensemble of similar systems, $m_i(\vec{x}, t) \equiv \langle M_i(\vec{x}, t) \rangle$.

One complete update requires the four substeps discussed above. Thus we consider the unit interval of time to be the time required to complete these four substeps. The result at $t' = t + 1/2$ (of transport and mixing along the horizontal direction) is

$$\begin{aligned}
 m_0\left(\vec{x}, t + \frac{1}{2}\right) &= \frac{1}{2} [m_0(\vec{x} - \hat{x}, t) + m_1(\vec{x} - \hat{x}, t)] \\
 m_1\left(\vec{x}, t + \frac{1}{2}\right) &= \frac{1}{2} [m_0(\vec{x} + \hat{x}, t) + m_1(\vec{x} + \hat{x}, t)].
 \end{aligned}
 \tag{13}$$

Now considering transport along the \hat{y} direction we find that

$$\begin{aligned}
 m_0(\vec{x}, t + 1) &= \frac{1}{2} \left[m_0(\vec{x} - \hat{y}, t + \frac{1}{2}) + m_1(\vec{x} - \hat{y}, t + \frac{1}{2}) \right] \\
 &= \frac{1}{4} [m_0(\vec{x} - \hat{x} - \hat{y}, t) + m_1(\vec{x} - \hat{x} - \hat{y}, t) + m_0(\vec{x} + \hat{x} - \hat{y}, t) + m_1(\vec{x} + \hat{x} - \hat{y}, t)].
 \end{aligned}
 \tag{14}$$

⁵ Note that when dealing with diffusion on a two-dimensional triangular lattice we would apply a sequence of three substeps, using the same two channels to transport data subsequently along each of the three principle lattice directions. Generalization of transport on a more complex higher-dimensional lattice would require addition of a substep along each new principle lattice direction.

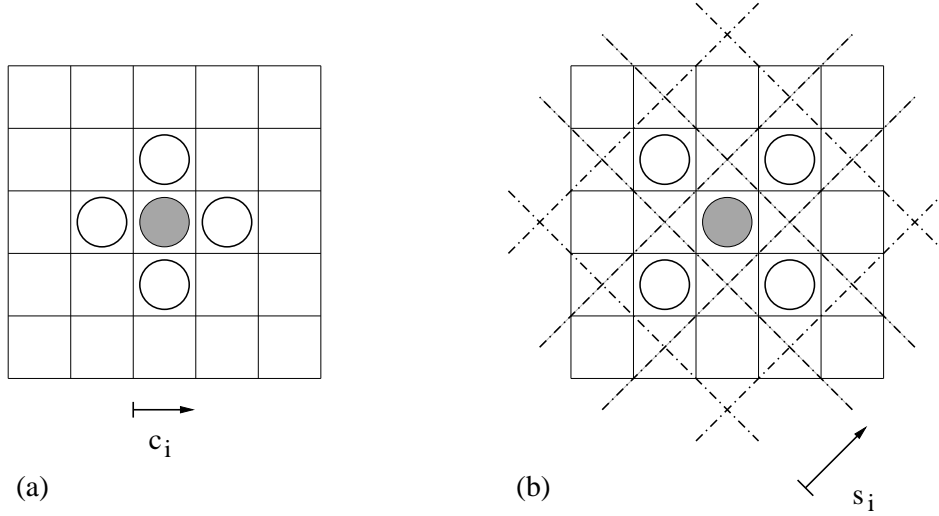


FIG. 3: A snapshot of the lattice. A single particle is shown as the shaded circle. The potential sites that a particle can occupy at the next iteration are shown as open circles. (a) Updating using the standard LGA diffusion algorithm. Note the lattice spacing, $|c_i|$. (b) Updating using the dimension-splitting technique. At the next iteration, the particle advances one step along an oblique lattice. This oblique lattice is shown superimposed on the original lattice. Note the lattice spacing, $|s_i| = \sqrt{2}|c_i|$.

Analogously,

$$m_1(\vec{x}, t+1) = \frac{1}{4} [m_0(\vec{x} + \hat{x} + \hat{y}, t) + m_1(\vec{x} + \hat{x} + \hat{y}, t) + m_0(\vec{x} - \hat{x} + \hat{y}, t) + m_1(\vec{x} - \hat{x} + \hat{y}, t)]. \quad (15)$$

Let us redefine the lattice vectors such that $\hat{s}_0 = \hat{x} + \hat{y}$, $\hat{s}_1 = -\hat{x} + \hat{y}$, $\hat{s}_2 = -\hat{s}_0$, and $\hat{s}_3 = -\hat{s}_1$. Note that the unit vectors of the oblique lattice are longer than those on the original lattice: $|\hat{s}_i| = \sqrt{2}|\hat{c}_i|$. The evolution of the density

$$\begin{aligned} \rho(\vec{x}, t+1) &= m_0(\vec{x}, t+1) + m_1(\vec{x}, t+1) \\ &= \frac{1}{4} \sum_{i=0}^3 \left[\sum_{j=0}^1 m_j(\vec{x} - \hat{s}_i, t) \right] = \frac{1}{4} \sum_{i=0}^3 \rho(\vec{x} - \hat{s}_i, t). \end{aligned} \quad (16)$$

Comparing this equation to Eq. (7), we see that the dimension-splitting model maps onto the conventional model since the oblique lattice vectors, \hat{s}_i , map onto the original lattice vectors, \hat{c}_i . Reference [23] and Ref. [24] contain many quantitative plots of the resulting diffusion. Instead of reproducing those plots here, we refer the reader to the original works.

2. Diffusion coefficients

We can easily control the coefficient of self-diffusion (D in Eq. (10)), a property which will prove useful in the models discussed in Sec. III. In particular, we can increase the diffusion constant by a factor of k^2 in one of two ways. First we can transport the particles by k lattice unit vectors each streaming step. Secondly, recalling that LGA models have transport steps alternating with interaction steps, we can run k^2 diffusion steps for each particle interaction step. With more effort, we can use a non-uniform $r(\vec{x}, t)$ as shown by Chopard and Droz[23, 24] to increase or decrease the probability of reversing direction at each interaction step and thereby decrease or increase the mean free path, and hence the diffusion constant.

C. Constraints and conserved quantities

As discussed above, the particles execute simultaneous random walks with exclusion. If we start with at most one particle per channel, that constraint is preserved by the dynamics. Moreover, the total number of particles—

alternatively viewed as mass, kinetic energy, or total energy—is conserved. Momentum is not conserved because particles can change their states of motion without a compensating change elsewhere.

Our model also has *spurious* conserved quantities (quantities that we do not necessarily wish to conserve but are conserved nonetheless). The dynamics partitions the lattice into four interleaved sublattices, and the number of particles on each sublattice is independently conserved. Consider a particle initialized at the center of the lattice (marked by the shaded circle) as shown in Fig. 3(b). At the next time step we know it will occupy one of the sites marked with an open circle. Likewise, only a particle initialized on a site marked by an open circle will be able to occupy the site marked with the shaded circle at the next time step. If we think of superimposing a checkerboard on the lattice, where we color the squares on one checkerboard sublattice “red” and those on the other “black”, it is clear that all particles on the red sites will occupy red sites at the end of any complete update of the system. Yet we can further divide each checkerboard into two. Again consider the situation illustrated in Fig. 3(b). A particle initialized on the shaded site cannot occupy that shaded site at the next time step (even though that shaded site is on the appropriate checkerboard). Thus only half of the sites on that checkerboard are accessible in one time step, and so we are simulating diffusion independently on each of four interleaved sublattices. In contrast, the Chopard-Droz model partitions the lattice into two checkerboard sublattices, simulating diffusion independently on each.⁶ If we increase the dimensionality of our model, we double the number of independent checkerboards with each additional dimension.

In Fig. 1, we started all four sublattices in the identical state, so the separate particle conservation on interleaved sublattices is not apparent. If we had started with one of the sublattices occupied, at any point in the dynamics only one of the sublattices would be occupied. When diffusion is used as a component of a larger simulation, these conservation laws are often broken by interactions between subsystems. For example, as discussed below, we consider interactions of diffusing vapor particles and a stationary solid, where the interactions mix the sublattices. A vapor particle will aggregate regardless of the sublattice ordering, and subsequent evaporation from the solid returns that vapor particle placed on an arbitrary sublattice. We can also eliminate the spurious conservation laws in a more intrinsic way: we simply simulate a single sublattice and embed it into the original lattice by shrinking it by a factor of two in each direction. Odd and even timesteps must be embedded with slightly different offsets (amounting to one-half the lattice spacing), but this is just a slight complication in the implementation. For a detailed discussion of the embedding technique see Ref. [32].

Often, when developing microscopic models, we wish to create models which are microscopically invertible (*i.e.*, models which conserve microscopic information). If we use an invertible random number generator for our source of random bits, we can build an invertible model of diffusion, since we can then undo the mixing permutation. (In the first example below, we use an independent invertible LGA as the generator of a dynamical random number field). We can always undo the transport steps (by transporting in the opposite direction). Thus both the Toffoli/Chopard-Droz model and our model can be made invertible.

III. APPLICATIONS

A. The reversible aggregation model

We developed this technique for simulating diffusion while building a model of cluster growth via the aggregation of particles which were previously diffusing along the lattice. This model is an extension of the athermal diffusion-limited aggregation (DLA) model[33]. We place a DLA lattice system in contact with a simulated heat bath and allow only microscopically reversible heat exchanges between the two subsystems.

During each aggregation event, a diffusing gas particle becomes a stationary cluster member while latent heat is released. The heat is explicitly modeled by introducing a diffusing heat particle onto a heat bath lattice. The inverse process is also allowed: a cluster member can absorb a heat particle and evaporate, becoming a gas particle. Initially the heat bath is empty, and we observe rapid, nonequilibrium growth. During this initial phase the cluster structures resemble those typical to DLA systems, such as the bushy clusters formed as frost collects on a window pane. During the subsequent slow approach to thermodynamic equilibrium, the cluster structures anneal until reaching the highest entropy macrostate allowed for a connected cluster in a finite volume: a branched polymer. For more details of the reversible aggregation (RA) model see Ref. [34].

⁶ The original lattice-diffusion model introduced by Toffoli uses a partitioning scheme which doesn’t split the space into uncoupled sublattices.




FIG. 4: (a) A cluster grown via a two-dimensional RA implementation. (b) A visualization of a three-dimensional RA cluster. The rendering is done using an LGA algorithm which simulates discrete light particles moving on the lattice[7]. Diffusive phenomena are at the foundation of the RA model: they are handled via the dimension-splitting technique. Note that extending the two-dimensional model to a three-dimensional model required merely adding one line of code.

This local, microscopic, reversible approach to modeling allows us to explicitly track how information flow at the microscopic scale gives rise to dissipation and properties such as increasing temperature at the macroscopic scale. We demonstrate that even far from thermodynamic equilibrium, as the simulated clusters are growing and annealing, we have a model with a well defined temperature, and that these models can serve as a numerical laboratory for investigating nonequilibrium thermodynamics. Moreover, discrete reversible systems such as this—which include local heat flow, creation of entropy, and macroscopic dissipation—seem to capture the full essence of the phenomena of pattern formation more completely than do irreversible models. For more details see Refs. [15, 34].

Simulations of the RA model were implemented and run on the CAM-8 lattice-gas computer[7]. Figure 4 shows a typical RA cluster, first for a two-dimensional implementation, then for a three-dimensional one. Implementing the three-dimensional model was a simple matter of including one more line of code in our program (for transport along the third dimension). Note that the diffusion technique forms the foundation of our model, which involves two interacting diffusion fields (a field of gas particles and a field of heat particles). We can control the diffusion coefficient independently for each field, as discussed in Sec. II B 2, and observe the change in the cluster structures and time to reach equilibrium as a function of this parameter.

B. A model for sedimentation

Another application which makes use of the dimension-splitting technique in a reaction-diffusion context is the Biofilm CA model. This is a model developed by Pizarro, Griffeath and Noguera[35, 36], and programmed by Shalizi and Margolus on the CAM-8[7].

The Biofilm CA simulates the growth of biological films (e.g. sludge, dental plaque) by having food and bacteria react and diffuse in two or more dimensions. The diffusion rate for the food depends on whether there is bacteria present, and the model includes birth, food sources, eating, death by starvation, random death, boundaries, attachment to a substrate, attachment to an aggregate, death of unattached cells, and the periodic removal of entire bacterial clusters which have become disconnected from the substrate (this simulates the action of water currents).

The diffusion in this model is handled one dimension at a time. Each diffusing species can occupy one or both (or neither) of two channels at each lattice site, and the contents of the channels are alternately swapped (or not swapped) at random, and then transported oppositely along one dimension. For the food particles, each channel contains a single bit (present or not). For the bacteria, each channel contains two bits (present or not, and if present how hungry). In regions which should have a longer mean free path, the mixing probability is reduced. Food diffuses in from the top of the space, bacterial growth starts on the substrate at the bottom. To reduce the time taken for food to initially reach the bacteria by diffusion, food starts off uniformly distributed across the space. For more details on the model see Ref. [36].




FIG. 5: Dimension-splitting biofilm simulation. (a) A two-dimensional simulation run on a 256×256 space. Food enters at the top, and a biofilm grows on the substrate at the bottom. (b) The same simulation run in three dimensions on a $256 \times 256 \times 256$ lattice. Rendering is done as in Fig. 4(b). We’re looking at a downward angle at the 256×256 substrate, coated with biofilm growth. The light source is on the left. The food particles are not shown.

The model was initially developed as a two-dimensional system, which could be run quickly and easily, and conveniently visualized. The model was tested and debugged in two-dimensions until it behaved as expected. A typical state in the growth of a two-dimensional biofilm is shown in Fig. 5(a). The model was then generalized to three dimensions without adding any bits to each lattice site or defining any new rules to apply to each site. The sequence of operations needed to update a single dimension was simply iterated three times instead of two, with the third iteration involving transport of diffusing particles along the added third dimension. An appropriate three-dimensional initial state was prepared, and generic three-dimensional visualization routines were activated. The model immediately behaved as expected in three dimensions—no further debugging was needed. Only fine tuning of initial densities, diffusion rates, etc., was required. A typical state in the growth of a three-dimensional biofilm is shown in Fig. 5(b).

IV. DISCUSSION

We have introduced a dimension-splitting technique for use in LGA models with diffusive particle transport. For an LGA of arbitrary dimensionality, interactions involving only one dimension at a time are sufficient for iterating the dynamics. Thus, a model can be developed in one dimension, and extended to more dimensions simply by adding additional fractional timesteps. The same dynamics are applied in each fractional timestep, with the same diffusing particles moving along each dimension in turn. By reinterpreting and reusing the same degrees of freedom and the same interactions at each fractional timestep, diffusion is performed with a small amount of state information.

Using this scheme, simple qualitative models involving diffusion can easily be constructed using a kind of stylized molecular dynamics. We can appeal to our intuition in constructing such models, simply including all relevant species and making them all interact in a physically reasonable manner. Constructing more realistic models involves more work, however, since the diffusion behavior of our model has only been demonstrated in the mesoscopic limit. To use the model only in this limit would involve running a very large number of diffusion steps for each particle interaction step. This might be resolved by requiring a finer lattice for the diffusers than for other entities in the simulation. For many models, however, we don’t expect that a quantitatively realistic interaction requires explicit mesoscopic averaging of the microscopic constituent simulations. Demonstrating realism for a given kind of simulation requires both theoretical work and experimental verification. Even models which are not completely realistic, however, may be viewed as “toy” physical systems which can be simulated exactly and analyzed theoretically[34].

It seems that there is an important cost for our model’s simplicity: the dynamic range in particle density at each lattice site is small. N bits of state at each site yield a maximum density of only N diffusers per site. Numerical simulations of diffusion, in contrast, allow up to 2^N diffusers to be represented at each site using the same amount of state, with only a linear increase in the computational complexity of the dynamics. However, this contrast is only apparent. The different species present at a lattice site can of course be given different power-of-two weights: each “1”

of species A might represent a single particle, each “1” of species B , two particles, etc. Independent dimension-split diffusion of each such species provides a large dynamic range while retaining exact particle conservation.

LGA models with as few as two bits of diffusing state at each lattice site can be constructed using the dimension-splitting technique. The benefits of similar economies in simulation state have recently been discussed in other contexts. For example, in the context of reversible compilers see Ref. [37]; in the context of systems relying on pseudorandom numbers see Ref. [38].

It might seem that the operation on one dimension at a time would result in slower simulations than a more complex LGA model, which mixes data in several dimensions at once. However, from a computational complexity standpoint, this is not at all obvious, since the more complex LGA model would have more data to interact. If, for example, diffusion is being handled by circuitry capable of complex operations, then using only simple operations may waste resources. If, on the other hand, only very simple operations are available in the circuitry, and only one operation can be performed at a time, then more data would require more operations and mean a slower simulation. If it is ever possible to perform simulations with molecular computers, the latter of these two situations seems much more likely than the former.

Our dimension-splitting diffusion technique is particularly well suited for use in concert with other dimension-splittable primitives. Composite models based only on such primitives share the properties of reduced state and the possibility of developing models in one or two dimensions, and then running them without change in higher dimensions.

A useful primitive of this type is the LGA-fluid attraction and repulsion technique used by Yepez and Appert in their momentum-conserving lattice-gas aggregation models[39, 40]. In these models, pairs of particles lying along principle lattice directions and separated by a given distance are brought together to interact. By using a set of distances and specified types of interactions (repulsive interactions turn both particles away from their midpoint, attractive interactions turn them towards it), prescribed potentials can be synthesized operating between particle aggregates.⁷

Since these models aren’t hydrodynamic, momentum conserving gas interactions used to mix particles going in different directions can also be made to interact one dimension at a time, forming a rest particle from each head-on collision which then decays into two new particles moving out in opposite directions—when the particle decays can be determined stochastically. Using the technique of the present paper, diffusing components can also be included in such models. Note that in all of these models, the fractional timesteps are applied along each of the principle lattice directions, so that in a two-dimensional triangular lattice model, for example, there are three transport steps rather than just two. Thus a better name for the technique might be “lattice-splitting” rather than “dimension-splitting”. It would be interesting to find additional examples of such splitting techniques which can be used as components of LGA models.

-
- [1] E. S. Boek, P. V. Coveney, H.N.W. Lekkerkerker, and P. van der Schoot. Simulating the rheology of dense colloidal suspensions using dissipative particle dynamics. *Phys. Rev. E*, 55:3124–3133, 1997.
 - [2] B. M. Boghosian, F. J. Alexander, and P. V. Coveney (Eds.). Proceedings of the sixth international conference on discrete models for fluid mechanics. *Int. J. Mod. Phys. C*, 8(4), 1997.
 - [3] G. H. Weiss. *Aspects and Applications of the Random Walk*. North-Holland, Amsterdam, 1994.
 - [4] J. Hardy, O. de Pazzis, and Y. Pomeau. Molecular dynamics of a classical lattice gas: transport properties and time correlation functions. *Phys. Rev. A*, 13(5):1949–1961, 1976.
 - [5] U. Frisch, B. Hasslacher, and Y. Pomeau. Lattice-gas automata for the Navier-Stokes equation. *Phys. Rev. Lett.*, 56(14):1505–1508, 1986.
 - [6] N. Margolus, T. Toffoli, and G. Vichniac. Cellular-automata supercomputers for fluid-dynamics modeling. *Phys. Rev. Lett.*, 56(16):1694–1696, 1986.
 - [7] N. H. Margolus. CAM-8: A computer architecture based on cellular automata. In A. Lawniczak and R. Kapral, editors, *Pattern Formation and Lattice-Gas Automata*. American Mathematical Society, 1996.
 - [8] D. H. Rothman and S. Zaleski. *Lattice-gas cellular automata : simple models of complex hydrodynamics*. Cambridge University Press, Cambridge, UK, 1997.
 - [9] J. P. Rivet and J. P. Boon. *Lattice Gas Hydrodynamics*. Cambridge Univ Press, Cambridge, UK, 2000.
 - [10] D. A. Wolf-Gladrow. *Lattice-Gas Cellular Automata and Lattice Boltzmann Models : An Introduction*. Lecture Notes in Mathematics. Springer-Verlag, New York, 2000.

⁷ Their original models didn’t use dimension-splitting but instead used “virtual particles” moving in all directions at once to carry the forces. In translating their models to the CAM-8 machine, we employed dimension-splitting.

- [11] R. Kapral, A. Lawniczak, and P. Masiar. Oscillations and waves in a reactive lattice-gas automaton. *Phys. Rev. Lett.*, 66(19):2539–2542, 1991.
- [12] J. P. Boon, D. Dab, R. Kapral, and A. Lawniczak. Lattice gas automata for reactive systems. *Physics Reports*, 273(2):55–147, 1996.
- [13] T. Toffoli and N. H. Margolus. *Cellular Automata Machines: A New Environment for Modeling*. MIT Press, Cambridge, MA, 1987.
- [14] M. A. Smith. *Cellular Automata Methods in Mathematical Physics*. PhD thesis, Massachusetts Institute of Technology, May 1994.
- [15] R. M. D’Souza. *Macroscopic order from reversible and stochastic lattice growth models*. PhD thesis, Massachusetts Institute of Technology, August 1999.
- [16] N. H. Margolus. Crystalline computation. In A. Hey, editor, *Feynman and Computation*. Addison-Wesley, 1998.
- [17] W. F. Ames. *Numerical methods for partial differential equations*. Academic Press, New York, 2nd edition, 1977.
- [18] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: the art of scientific computation*. Cambridge University Press, Cambridge, 2nd edition, 1992.
- [19] T. Toffoli. Three-dimensional rotations by three shears. *Graphical Models and Image Processing*, 59:89–96, 1997.
- [20] To our knowledge, the use of alternating directions in cellular automata models was first proposed by Edward Fredkin in the context of modeling fundamental physics. Personal communication, circa. 1990.
- [21] O. L. Bandman. Comparative study of cellular-automata diffusion models. In *Lecture notes in computer science*, volume 1662, pages 395–409, 1999.
- [22] P. V. Coveney, J. B. Maillet, J. L. Wilson, P. W. Fowler, O. Al-Mushadani, and B. M. Boghosian. Lattice-gas simulations of ternary amphiphilic fluid flow in porous media. *Int. J. of Mod. Phys. C*, 9(8):1479–1490, 1998.
- [23] B. Chopard and M. Droz. Cellular automata model for the diffusion equation. *J. Stat. Phys.*, 64(3/4):859–892, 1991.
- [24] B. Chopard and M. Droz. *Cellular Automata Modeling of Physical Systems*. Cambridge University Press, Cambridge, UK, 1998.
- [25] R. Byron Bird, E. N. Lightfoot, and W. E. Stewart. *Transport Phenomena*. John Wiley & Sons, New York, 2nd edition, 2001.
- [26] D. Walgraef. *Spatio-Temporal Pattern Formation : With Examples from Physics, Chemistry, and Materials Science*. Springer Verlag, New York, 1997.
- [27] H. C. Berg. *Random Walks in Biology*. Princeton University Press, Princeton, 1983.
- [28] J. Crank. *The Mathematics of Diffusion*. Oxford University Press, Oxford, 2nd edition, 1975.
- [29] H. D. Chen, B. M. Boghosian, P. V. Coveney, and M. Nekovee. A ternary lattice Boltzmann model for amphiphilic fluids. *Proc. Roy. Soc. of London, Series A*, 456:2043–2057, 2000.
- [30] J. L. Lebowitz. Microscopic reversibility and macroscopic behavior: Physical explanations and mathematical derivations. In J. J. Brey, J. Marro, J. M. Rubi, and M. San Miguel, editors, *Lecture Notes in Physics*. Springer, 1995.
- [31] D. Stauffer. Irreversible statistical mechanics from reversible motion: Q2R automata. *Comp. Phys. Comm.*, 127(1):113–119, 2000.
- [32] N. Margolus. An embedded DRAM architecture for large-scale spatial-lattice computations. In *The 27th Annual International Symposium on Computer Architecture*. IEEE Computer Society, 149–160, (2000).
- [33] T. A. Witten and L. M. Sander. Diffusion-Limited Aggregation, a kinetic critical phenomenon. *Phys. Rev. Lett.*, 47(19):1400–1403, 1981.
- [34] R. M. D’Souza and N. H. Margolus. Thermodynamically reversible generalization of diffusion limited aggregation. *Phys. Rev. E.*, 60(1):264–274, 1999.
- [35] G. E. Pizarro, D. Griffeath, and D. R. Noguera. Quantitative cellular automaton model for biofilms. *J. Environ. Eng.*, 2001. In press.
- [36] G. E. Pizarro. *Quantitative modeling of heterogeneous biofilms using cellular automata*. PhD thesis, University of Wisconsin, 2001.
- [37] C. Carothers, K. Perumalla, and R. Fujimoto. Efficient optimistic parallel simulations using reverse computation. In *ACM Transactions on Modeling and Computer Simulation*, volume 9, 1999.
- [38] N. Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.
- [39] J. Yepez. Lattice-gas crystallization. *J. Stat. Phys.*, 81(1/2):255–294, 1994.
- [40] C. Appert and S. Zaleski. Lattice gas with a liquid-gas transition. *Phys. Rev. Lett.*, 64, 1990.